arXiv:1612.00907v1 [math.NA] 3 Dec 2016

Multigrid In Energy Preconditioner for Krylov Solvers

R. N. Slaybaugh, T. M. Evans[a], G. G. Davidson[a], P. P. H. Wilson[b]

Department of Mechanical Engineering and Material Science, University of Pittsburgh, 605 Benedum Hall, 3700 O'Hara Street, Pittsburgh, PA 15261, USA

[a]*Radiation Transport Group, Oak Ridge National Laboratory, PO BOX 2008 MS6170, Oak Ridge TN 37831, USA*

[b]*Department of Nuclear Engineering and Engineering Physics, University of Wisconsin - Madison, 419 ERB, 1500 Engineering Drive, Madison, WI 52706, USA*

Abstract

We have added a new multigrid in energy (MGE) preconditioner to the Denovo discrete-ordinates radiation transport code. This preconditioner takes advantage of a new multilevel parallel decomposition. A multigroup Krylov subspace iterative solver that is decomposed in energy as well as space-angle forms the backbone of the transport solves in Denovo. The space-angle-energy decomposition facilitates scaling to hundreds of thousands of cores. The multigrid in energy preconditioner scales well in the energy dimension and significantly reduces the number of Krylov iterations required for convergence. This preconditioner is well-suited for use with advanced eigenvalue solvers such as Rayleigh Quotient Iteration and Arnoldi.

Keywords: Preconditioning, Multigrid, Krylov, Neutron Transport

1. Introduction

The steady-state Boltzmann equation for neutron transport finds "where all the neutrons are" in a nuclear system. The more accurately this is known, the more accurately new systems can be developed. Challenging transport problems today are three-dimensional with up to thousands $\times$ thousands $\times$ thousands of mesh points, use up to $\sim$150 energy groups, include accurate expansions of scattering terms, and are solved over many angular directions. Such discretizations create systems with 10^{9-10} coupled algebraic equations. High-fidelity, coupled,

Email address: rns37@pitt.edu (R. N. Slaybaugh)

multi-physics calculations are the next generation of "grand challenge" problems for reactor analysis, requiring that the finely-resolved neutron flux be calculated quickly and accurately.

Very large computers, like the the Jaguar machine [1], are now available to perform such high-fidelity calculations. Access to such machines has changed the types of neutron transport problems that can be solved practically. However, most existing solution methods are not able to take full advantage of new computer architectures, or they have convergence properties that limit their usefulness for difficult problems.

The deterministic transport equation is often solved with iterative methods, and Krylov methods are becoming an increasingly popular choice. Low-cost preconditioners that can reduce the number of iterations needed for convergence are invaluable. Benzi's 2002 survey paper on preconditioning techniques for large linear systems states "it is widely recognized that preconditioning is the most critical ingredient in the development of efficient solvers for challenging problems in scientific computation"[2]. Ten years later, this is still true. The need for preconditioning is particularly high for Krylov methods because the memory required and cost per iteration increase dramatically with the number of iterations [2].

This work aims to improve convergence behavior for 3-D transport calculations with a new preconditioner that can use leadership-class machines fully. A multigrid in energy preconditioner (MGE) that uses a Krylov solver to do the multigroup transport solves has been added to the code Denovo [3]. The MGE preconditioner reduces the number of Krylov iterations in both fixed source and eigenvalue problems. This success has been demonstrated on large, realistic problems of interest. The preconditioner is particularly well-suited for the fine energy structures needed in challenging discrete-ordinates calculations.

This paper is organized as follows: Section 2 gives background information about solving the transport equation. Section 3 provides an overview of multigrid preconditioners used in the nuclear community. The MGE method, usage options, and how it is parallelized are presented in Section 4. Finally, test results demonstrating the preconditioner's impact and the associated conclusions are presented in Sections 5 and 6, respectively.

2. Background

To understand MGE and how it can enable the use of leadership-class computers, the mathematical details of the transport equation and how it is solved must be discussed first.

2.1. Transport Equation

The fixed source form of the steady state Boltzmann transport equation is

$$[\hat{\Omega} \cdot \nabla + \Sigma(\vec{r}, E)]\psi(\vec{r}, \hat{\Omega}, E) = q_{ex}(\vec{r}, \hat{\Omega}, E) + \int_0^\infty dE' \int_{4\pi} d\hat{\Omega}' \, \Sigma_s(\vec{r}, E' \to E, \hat{\Omega}' \cdot \hat{\Omega})\psi(\vec{r}, \hat{\Omega}', E') \, , \quad (1)$$

where the quantities are at location $\vec{r}$, at energy E, and are traveling in direction $\hat{\Omega}$. The angular neutron flux in neutrons per unit length squared per steradian is $\psi(\vec{r}, \hat{\Omega}, E)$, which expresses the location of the neutrons in phase space; q_{ex} is an external source; Σ_x is the cross section that represents the likelihood of interaction x (here s is for scattering, f is for fission, and no subscript indicates all interactions) in units of inverse length [4].

In the presence of fission, the external source is replaced by a fission source,

$$\chi(E) \int_0^\infty dE' \, \nu\Sigma_f(\vec{r}, E') \int_{4\pi} d\hat{\Omega}' \, \psi(\vec{r}, \hat{\Omega}', E') \, . \quad (2)$$

Here $\chi(E)$ is the spectrum specifying the energy distribution of neutrons born from fission, and ν is the average number of neutrons released per fission [4].

It is often of interest to know the asymptotic behavior of a system with fission. A reactor is called "critical" if the chain reaction is self-sustaining and time-independent. If the system is not in equilibrium, then the asymptotic neutron distribution, or fundamental mode, will grow or decay exponentially over time. A convenient way to capture this behavior is to assume ν can be adjusted to obtain a time-independent solution by replacing it with $\frac{\nu}{k}$, where k is the parameter expressing the deviation from critical. A spectrum of eigenvalues can be found, but at long times only the non-negative solution corresponding to the largest real eigenvalue will dominate [4].

A brief aside about some properties of the transport equation will aid in understanding the solution techniques. In a void, the transport equation is like a hyperbolic wave equation. For highly-scattering regions where Σ_s is close to Σ, the equation becomes elliptic for the steady-state case. If the scattering is forward-peaked then the equation is parabolic. All of these classes of equations have different solution strategies, so it is difficult to find one method that works well in all cases of physical interest [5].

To numerically solve Equation (1) it is discretized in space, angle, and energy. This work uses the multigroup energy approximation, the scattering term is expanded in Legendre polynomials (P_l), and discrete ordinates (S_N) are used

to treat direction of neutron travel. There are many spatial differencing methods available, the discussion of which is beyond the scope of this document as the proposed work is not dependent upon the spatial discretization employed [6]. To ensure the new methods apply to the most general cases, it will be assumed that the matrices resulting from discretization are not necessarily symmetric.

After all of the discretizations are performed, the multigroup S_N equations can be written in operator form as

$$\mathbf{L}\psi = \mathbf{MS}\phi + q \,, \qquad \text{(fixed source)} \tag{3}$$

$$\mathbf{L}\psi = \mathbf{MS}\phi + \frac{1}{k}\mathbf{M}\chi f^T\phi \,. \qquad \text{(eigenvalue)} \tag{4}$$

Here $\mathbf{L}$ is the first-order linear differential transport operator; $\mathbf{M}$ is the moment-to-discrete operator that projects the angular flux moments, ϕ, onto discrete angles; $\mathbf{S}$ is the scattering matrix; f contains the fission source, $\nu\Sigma_f$; and q is a source term. The angular flux moments are related to the angular flux through the discrete-to-moment operator:

$$\phi = \mathbf{D}\psi \,. \tag{5}$$

Using this relationship, Equations (3) and (4) can be rearranged such that they are a function of only ϕ. The formulation is aided by defining $\mathbf{T} = \mathbf{DL}^{-1}$ and $\mathbf{F} = \chi f^T$ [7]:

$$(\mathbf{I} - \mathbf{TMS})\phi = q \,, \qquad \text{(fixed source)} \tag{6}$$

$$(\mathbf{I} - \mathbf{TMS})\phi = \frac{1}{k}\mathbf{TMF}\phi \,. \quad \text{(eigenvalue)} \tag{7}$$

Once the matrices are multiplied together, a series of single "within-group" equations that are each only a function of space and angle result. If the groups are coupled together by neutrons scattering from a low energy group to a higher energy group, then iterative "multigroup" solves over the coupled portion of the energy range may be required. If the eigenvalue is desired, an additional "eigenvalue" solve is needed [8].

2.2. Multigroup Solve

Traditionally, the multigroup solve has been done with Gauss Seidel (GS). GS is iterative in energy. A space-angle solve using a within-group solver, such as source iteration or a Krylov method, is performed for each energy group in series.

The groups are solved from $g = 0$, the highest energy, to $g = G$, the lowest. For a group g and an energy iteration index j this is [3]

$$\big(\mathbf{I} - \mathbf{TMS}_{gg}\big)\phi_g^{j+1} = \mathbf{TM}\big(\sum_{g'=0}^{g-1} \mathbf{S}_{gg'}\phi_{g'}^{j+1} + \sum_{g'=g+1}^{G} \mathbf{S}_{gg'}\phi_{g'}^{j} + q_g\big)\,. \tag{8}$$

The first term on the right includes downscattering contributions from higher energies, and the second term represents upscattering contributions from lower energy groups that have not yet been converged for this energy iteration. Groups that only contain downscattering are simply solved once since the second term on the right is zero. Groups with upscattering, however, must be iterated until they converge. Convergence of GS is governed by the spectral radius of the system, so the method can be very slow when upscattering has a large influence on the solution [5]. GS is fundamentally serial in energy because of how the group-to-group coupling is treated.

A recently-added multigroup (MG) Krylov solver removes the traditional "within-group" / "multigroup" iteration structure. This allows the solver to handle upscattering efficiently, and enables parallelization in the energy dimension. The solver has been shown to successfully scale to hundreds of thousands of cores. For example, a fixed-source test scaled from 69,102 cores to 190,080 cores with 98% efficiency [9].

The MG Krylov solver combines the space-angle and energy iterations to make one space-angle-energy iteration level. This allows the energy groups to be decomposed such that they can be solved in parallel. The space-angle-energy iterations are much like the within-group space-angle iterations, except that the iteration is over a block of groups instead of just one group. The block can include all groups or just upscattering groups, in which case the downscattering groups are treated in series as before. An added benefit of this solver is that Krylov methods generally converge more quickly than GS [10].

The multigroup Krylov method applied to the upscattering block is shown here, where $\mathbf{S}_{\text{up_block}}$ contains the upscattering groups and $\mathbf{S}_{\text{up_source}}$ has the downscattering only groups:

$$\underbrace{\big(\mathbf{I} - \mathbf{TMS}_{\text{up_block}}\big)}_{\tilde{\mathbf{A}}}\phi_{\text{up_block}}^{n+1} = \mathbf{TM}(\mathbf{S}_{\text{up_source}}\phi_{\text{up_source}}^{n+1} + q)\,. \tag{9}$$

Trilinos [11] provides Denovo's Krylov solver, with a choice of either GMRES or BiCGSTAB [3]. The Krylov solver is given an operator that implements the

action of $\tilde{\mathbf{A}}$, or the matrix-vector multiply and sweep. In the MG Krylov solver, $\tilde{\mathbf{A}}$ is applied to an iteration vector, v, containing the entire upscattering block instead of just one group:

1. matrix-vector multiply: $y = \mathbf{M}\mathbf{S}_{\text{up_block}}v$,
2. sweep: $z = \mathbf{T}y$,
3. return: $v \leftarrow v - z$.

To implement the energy parallelization, the problem is divided into energy sets, with groups distributed evenly among sets. After each set performs its part of the matrix-vector multiply, a global reduce-plus-scatter is the only required inter-set communication. Since each set uses the entire spatial mesh with the same spatial decomposition, the established performance of spatial scaling does not change. The space-angle decomposition in Denovo comes from the KBA wavefront algorithm [12]. Evans et al. [6] showed scaling in space is limited to about 20,000 cores with a 500 million cell spatial mesh.

The added energy decomposition offers the ability to further decompose a problem, even if the performance limit of spatial decomposition has been reached. The total number of cores is equal to the number of computational domains, that is, the product of the number of energy sets and the number of spatial blocks. For 20,000 spatial blocks and 10 energy sets, which is a reasonable decomposition, 200,000 cores will be used. See Ref. [7] for more details.

The MG Krylov solver makes it possible to use very large machines and fine energy structures. The addition of this solver is one of the key motivators for the MGE preconditioner because Krylov methods often need to be preconditioned to converge in a small number of iterations. A preconditioner that can scale well and handle many energy groups is essential for using the MG Krylov solver.

2.3. Eigenvalue Solve

When the eigenvalue form of the transport equation is used, another iteration level is needed outside of the multigroup solves. A common way to solve the eigenvalue problem is with power iteration (PI). This method is attractive because it only requires matrix-vector products and two vectors of storage space. PI's convergence can be very slow for problems of interest, however, because the error is reduced in each iteration by a factor of $\mathbf{A}$'s dominance ratio, $\frac{\lambda_2}{\lambda_1}$, where the λs are $\mathbf{A}$'s eigenvalues ordered as $|\lambda_1| > |\lambda_2| \geq \cdots \geq |\lambda_n| \geq 0$.

For the standard form of power iteration, we define an energy-independent

eigenvalue, $\Gamma = f^T \phi$. Then, Eq. (7) can be written as

$$\begin{aligned} \mathbf{A}\Gamma &= k\Gamma \,, \qquad \text{where} \\ \mathbf{A} &= f^T (\mathbf{I} - \mathbf{TMS})^{-1} \mathbf{TM}\chi \,. \end{aligned} \tag{10}$$

Note that this is the ordinary form of the eigenvalue problem with eigenpair (Γ, k). The power iteration method, with i as the iteration index, is then

$$\Gamma^{i+1} = \frac{1}{k} \mathbf{A} \Gamma^i \,. \tag{11}$$

Inside of power iteration, the application of $\mathbf{A}$ to Γ requires the solution of a multigroup problem that looks like a fixed source problem,

$$(\mathbf{I} - \mathbf{TMS}) y^i = \mathbf{TM}\chi\Gamma^i \,, \tag{12}$$

$$\Gamma^{i+1} = f^T y^i \,. \tag{13}$$

Power iteration is just one of the eigenvalue solvers in Denovo. An Arnoldi solver and Rayleigh Quotient Iteration (RQI) are also available. For details about these methods and their implementation in Denovo see references [7] and [13], respectively.

3. Multigrid Preconditioning

For many kinds of problems the total number of multigroup iterations required to adequately converge the transport equation is large. To reduce iteration count, a preconditioner can be applied to transform the system of interest into another equivalent system that has more favorable properties and is easier to solve.

Right preconditioning leaves the right hand side of the equation unaffected and does not change the norm of the residual, which is used for convergence testing in most iterative methods. Right preconditioning is therefore often preferred over left or split preconditioning for iterative solvers [14]. A right preconditioner was implemented in this work, and the remaining discussion will be presented in right preconditioner format.

Let $\mathbf{G}$ be a non-singular preconditioner. Then $\mathbf{A}x = b$ can be transformed as

$$\mathbf{A}\mathbf{G}^{-1} y = b, \qquad x = \mathbf{G}^{-1} y \,. \tag{14}$$

The matrix $\mathbf{A}\mathbf{G}^{-1}$ is not formed in practice. The preconditioner can be applied by using some method to solve $\mathbf{G}y = c \rightarrow y \approx \mathbf{G}^{-1}c$, or by otherwise implementing the action of $\mathbf{G}^{-1}$ without explicitly forming and inverting $\mathbf{G}$ [2].

There are many different types of preconditioners. These methods often are dependent upon the choice of spatial discretization employed, only apply to within-group iterations, or have other important limitations. As a result, preconditioning Krylov methods for solving the 3-D neutron transport problem is an active and vital area of research. We elected to investigate multigrid preconditioners with the goal of making a preconditioner that is independent of spatial discretization and can take advantage of the MG Krylov solver.

Multigrid methods are stationary iterative schemes that can be used alone or as preconditioners for other methods. Beginning in the late 1980s the nuclear community started using spatial multigrid and/or angular multigrid as both solvers and preconditioners. The first use of spatial multigrid for transport equations in 1-D and 2-D was investigated by Nowak et al. Since that time, multigrid has been used in multiple dimensions, for both isotropic and anisotropic scattering, and for various spatial discretizations [5]. Some highlights from recent work are discussed below. All are applied to the S_N neutron transport equation unless otherwise noted.

In 1998, multigrid in space and multigrid in angle were used as preconditioners for Krylov methods and were tested for the 1-D, one-group, modified linear discontinuous (MLD) neutron transport equations by Oliveira and Deng [15]. They looked at isotropic scattering without absorption, isotropic scattering with absorption, and anisotropic cases. They had better results with multigrid than when using an Incomplete LU factorization method as a preconditioner.

In 2007, Chang et al. [16] used 2-D spatial multigrid for the isotropic scattering case, with corner balance finite difference in space and a four-color block-Jacobi relaxation scheme. A bilinear interpolation operator and its transpose were used for grid transfer. The method had some trouble with heterogeneous problems. The authors assert their algorithm is parallelizable.

In 2010, Lee [17] developed a method to do multigrid in space and angle simultaneously for two and three dimensions, isotropic and anisotropic scattering, one energy group, and a variety of spatial discretizations. The method can perform multigrid in only space, only angle, or some combination thereof. It also handles thick and thin cells.

3.1. Two-Grid

The two-grid acceleration method developed by Adams and Morel [18] was one of the earlier spatial multigrid methods, and it has been built upon by others. It is intended to accelerate convergence of the outer iterations when upscattering is present. The outer iteration method is Gauss Seidel and acceleration is only

applied to upscattering groups. The original work was done for slab geometries with a linear discontinuous (LD) discretization.

The two-grid method finds a correction for the isotropic component of the solution by solving a collapsed, one-group diffusion equation for the error. It is assumed that the zeroth moment of the error is a product of a spectral shape function and a space-dependent modulation function. To find the shape function, Fourier analysis is performed on GS. The zeroth moment of the cross sections is used to create the Fourier matrix, and then an eigenproblem is formed. Because the shape function is material-dependent, one such calculation is needed in each material region. Note that by using the zeroth moment, only the isotropic component of the solution is accelerated.

The diffusion equation in the two-grid approach differs slightly from the standard diffusion equation in that there is an extra term containing the gradient of the shape function. This is zero in homogeneous regions, but undefined at material interfaces. Adams and Morel found that neglecting the gradient term all together still gave good results for their test problems. This may not be true for more complex cases.

The two-grid method requires the diffusion operator to be consistent with the transport operator. This requirement can be difficult to meet for multi-dimensional problems, particularly for some spatial discretizations.

3.2. Two-Grid in Denovo

Denovo has a two-grid acceleration scheme based on the one developed by Adams and Morel. As mentioned, the iteration procedure uses a collapsed one-group diffusion equation to correct the low-order Fourier modes. Adams and Morel showed that the slowest converging spatial modes are diffusive and can be exactly computed in the infinite homogeneous case [18].

The original method can fail for systems of interest because of the consistency requirement for the discretization of the diffusion operator in multi-dimensional and multi-material problems. To make it applicable for the desired cases, the original two-grid method was modified by Evans et al. [6] by using a one-group transport equation instead of the diffusion equation. The modified method is called transport two-grid (TTG). The TTG method was derived to give the correct error estimation in the diffusive limit noted above.

To execute the TTG scheme in Denovo, a transport sweep is conducted in each group, a residual is calculated, the low-order transport solve is performed, an error form the transport sweep is calculated, and the scalar flux is updated. As

with the original two-grid method, the TTG method is limited to correcting only the isotropic flux moments.

The additional cost of the method is like solving one extra group, so for cases with many upscattering groups the cost can be amortized. The acceleration equation can be preconditioned with Diffusion Synthetic Acceleration for additional speed. Finally, a reduced quadrature set can be used inside TTG to limit the cost.

Some tests have shown TTG to be quite effective in improving the speed of convergence for upscatter problems when compared to unaccelerated GS [6]. While TTG solves the spatial discretization issue, it still only works with GS and the traditional within-group + multigroup iteration structure.

4. Multigrid In Energy

Preconditioning is important for increasing the robustness of Krylov methods. This is particularly true for the multigroup Krylov solver. This solver can create large Krylov subspaces because it forms the subspaces with multiple-group-sized vectors. As a result, any reduction in iteration count will have a significant benefit in terms of memory and cost per iteration. Prior to this work, there was no preconditioner in Denovo that could work with the MG Krylov solver.

The new preconditioner does multigrid in the energy dimension. To understand why multigrid in energy makes sense for neutron transport, some highlights about these methods are discussed here. See *A Multigrid Tutorial* by Briggs, Henson, and McCormick [19] for a thorough and approachable explanation of what multigrid methods are and why they work.

The error in x_i, the ith guess for $\mathbf{A}x_i = b_i$, can be written as a combination of Fourier modes. Each Fourier mode has a frequency, and the frequencies can range from low-frequency (smooth) to high-frequency (oscillatory). Iterative methods, also referred to as smoothers or relaxers, remove high-frequency error components quickly, but take many iterations to remove the low-frequency ones.

The idea of multigrid methods is to take advantage of the smoothing effects of iterative methods by making smooth errors look oscillatory and thus easier to remove. Errors that are low-frequency on a fine grid can be mapped onto a coarser grid where they are high-frequency. A relaxer is applied on the coarser grid to remove the now oscillatory error components. The remaining error is mapped to a still coarser grid and smoothed again. The problem is restricted to coarser and coarser grids until it is on a grid which is coarse enough to directly invert the equations.

Next, the coarsest result is prolonged back to the next-finer grid and used to correct the solution there. A few relaxations are done on this finer grid. The errors are prolonged back up the chain, continuously correcting on finer grids, until the finest grid is reached. This entire process is called a V-cycle.

Multigrid methods are differentiated by how many times the V-cycle is done and how many grids are used. The optimal combination of grids and cycles may depend on problem type. The addition of more grids and cycles will reduce error, but at an added cost.

A multigrid method was selected for many reasons. Multigrid methods remove the low-frequency error modes that require many Krylov iterations. The preconditioner was designed to take advantage of the energy decomposition used by the MG Krylov method. Each energy set can do work on its own grids and does not need to communicate with other energy sets. This is a communication savings compared to using grids in space or angle. An additional benefit is the simplicity of energy grids. Energy is one-dimensional, which allows for simpler coarsening and refinement than spatial or angular grids.

4.1. Method

Recall that right preconditioning is applied as $\mathbf{AG}^{-1}\mathbf{G}\phi = b$, where $\mathbf{A} = \mathbf{I} - \mathbf{TMS}$. To implement this in Denovo, y is defined as $\mathbf{G}\phi$ and the problem is broken into two steps:

1. with a Krylov method solve

$$\mathbf{AG}^{-1}y = b\,, \tag{15}$$

2. after finding y, calculate

$$\phi = \mathbf{G}^{-1}y\,. \tag{16}$$

To build energy grids, the energy group structure is coarsened so that each lower grid has fewer groups. The finest grid is the input energy structure, and the coarsest grid has one or a few groups. Each level has half as many groups as the previous level, rounded up if applicable. If there are $G+1$ groups on the fine grid there will be either $\frac{G+1}{2}$ or $\frac{G+2}{2}$ groups on the coarse grid. This is conceptually straightforward because the energy groups can be combined (restricted) and separated (prolonged) linearly.

There are a variety of options that must be considered in designing a multigrid scheme: the restriction and prolongation operators, the relaxation method, the number and/or pattern of V-cycles to use, the number of relaxations to do on

each grid level, and the depth of the V-cycle. Among these, the number of V-cycles, the number of relaxations per level, and the depth of the V-cycle have been implemented as user input options; the others are fixed.

The implemented restriction operator is a simple averaging scheme. Neighboring fine data are averaged together to make coarse data. Recall that in multigrid methods the variables being restricted and prolonged are error modes. For a grid with spacing h and a next-coarser grid $2h$, the errors are restricted as $e_g^{2h} = \frac{1}{2}(e_{2g}^h + e_{2g+1}^h)$ for $g = 0, ..., G_{end}$. If G is even then $G_{end} = \frac{G}{2}$; if G is odd then $G_{end} = \frac{G-1}{2}$ and the lowest energy group's datum is just copied. This scheme was chosen so that the thermal energy groups would be more resolved, which should improve accuracy for thermal reactors. The errors are restricted every time there is a transfer to a coarser grid.

The cross sections are restricted from the finest to the coarsest grid during problem initialization, and they do not change thereafter. The total and fission cross sections are restricted in the same way as the errors. Scattering is slightly more complicated since it has two indices, g and g'. When there are an even number of groups all cross sections are treated the same way: for $g' = 0, ..., G'_{end}$ and $g = 0, ..., G_{end}$,

$$\begin{aligned}\Sigma_s^{2h}(g, g') =& \frac{1}{4}[\Sigma_s^h(2g, 2g') + \Sigma_s^h(2g+1, 2g') \\ &+ \Sigma_s^h(2g, 2g'+1) + \Sigma_s^h(2g+1, 2g'+1)] \,. \end{aligned} \tag{17}$$

Unless there is upscattering in every group, some of the entries in Equation (17) will be zero. When there are an odd number of groups, the cross sections for the last g and g' groups are

$$\Sigma_s^{2h}(G_{end}, G'_{end}) = \Sigma_s^h(2G, 2G') \,, \tag{18}$$

for $g' = 0, ..., G'_{end}$:

$$\Sigma_s^{2h}(G_{end}, g') = \frac{1}{2}[\Sigma_s^h(2G, 2g') + \Sigma_s^h(2G, 2g'+1)] \,, \tag{19}$$

for $g = 0, ..., G_{end}$:

$$\Sigma_s^{2h}(g, G'_{end}) = \frac{1}{2}[\Sigma_s^h(2g, 2G') + \Sigma_s^h(2g+1, 2G')] \,. \tag{20}$$

The cross sections could be flux-weighted and re-restricted every time a new value for the flux were available, i.e. every new application of the preconditioner, to more accurately preserve physics. The experience of the computational community, however, has been that preconditioners do not need to rigorously and

accurately preserve physics to be effective. The trade-off between the physics preservation and the performance of the preconditioner cannot be known *a priori*. The assumption in this work is that the cost of recomputing the cross sections in every preconditioner application is greater than the benefit of more accurately representing the physics.

To prolong from a coarse to a fine grid, the points that line up between the grids are mapped directly: $e^h_{2g} = e^{2h}_g$ for $g = 0, ..., G_{end}$. To fill in the intermediate points on the fine grid, the adjacent coarse values are averaged: $e^h_{2g+1} = \frac{1}{2}(e^{2h}_g + e^{2h}_{g+1})$ for $g = 0, ..., G_{end} - 1$. In the e^h_{2g+1} case, $G_{end} = \frac{G-1}{2}$ when G is even to preserve the direct transfer of points that line up between the grids. The errors are prolonged every time there is a transfer to a finer grid. Since cross sections do not change, they are never prolonged.

There are other restriction and prolongation operators that are more rigorous and would preserve more accuracy when transferring between grids than those implemented. For example, a full weighting restriction operator would be more rigorous in combination with the current prolongation method than the current restriction operator is [19]. This change would be straightforward to implement.

An example of a more accurate and complex prolongation operator is to use shape functions to prolong from a coarse grid to a fine grid. The shape functions could be based on the previous iterate of the fine-grid vector, some known desirable expansion, etc. Such a change would be more difficult to implement and require more research.

These issues were not investigated in this work since the chosen operators, which were simple to code and check for correct implementation, were sufficient in practice. Any error added by the restriction and prolongation operators was largely removed by the relaxations. In addition, because this is a preconditioner, all of the pieces do not need to be rigorous. Complex grid transfer operators are unlikely to be worth the expense; however, it may be of value to investigate other simple operators.

The user chooses the number of V-cycles done for each preconditioner application. One V-cycle proceeds from the finest grid to the coarsest grid and back to the finest. The input option specifies the number of V-cycles that are concatenated together, with a default of 2. Each additional V-cycle should remove more error, but has a computational cost.

The depth of the V-cycle can also be specified by the user. The default behavior is determined by the number of groups, such that the grids will be coarsened until

there is only one energy group. The number of grids needed is given by [20]

$$\text{floor}\big(\log_2(G-1)\big)+2\,. \tag{21}$$

How this is handled when using energy sets is discussed below.

Some number of relaxations are performed on each level while traversing down and up the grids in a V-cycle. The number of relaxations per level is a user input choice with a default of 2. Performing more relaxations per grid should remove more error, but has a computational cost. The implemented relaxation method is weighted Richardson iteration. When applied to the transport equation, we obtain

$$\phi^m = \big(\mathbf{I}+\omega(\mathbf{TMS}-\mathbf{I})\big)\phi^{m-1}+\omega b^{m-1}\,, \tag{22}$$

where ω is a constant selected by the user that defaults to 1.

An important principle is that the preconditioner is only attempting to approximately invert $\mathbf{A}$. It is therefore reasonable to use a less accurate angular discretization in the preconditioner than the rest of the code. For example, the whole problem may be solved at S_{10}, but the preconditioner could only use S_2. There is an input option to specify an angular quadrature set to use in the preconditioner different from the angular quadrature set used in the rest of the problem; the default is to use the same set in both. At this time, this option has only been implemented for vacuum boundary conditions.

4.2. Parallelization

A key attribute of this preconditioner is that it is parallelizable in energy. This is because it uses the energy sets introduced by the MG Krylov solver. There are two ways to handle energy grids and energy sets together. One way is to restrict from G groups down to 1 group just as if there were no energy sets. This requires cross-set communication as soon as there are fewer groups than sets. This also causes some logistical difficulties related to what data are held by which sets at various points in the calculation.

The other approach is to prohibit cross-set communication by having each set do its own "mini" V-cycle. Each set restricts, prolongs, and relaxes on only its own groups. This strategy requires there to be at least two groups on every set. With an unequal number of groups per set, there is a choice between forcing all sets to have the same grid depth or allowing those with more groups to have deeper Vs. The first option enforces energy load balancing between sets, while the second allows the sets with more groups to remove more error.

The preconditioner is implemented so that all sets use the same grid depth. The benefit of load balancing is likely to be greater than having some sets use an extra grid. Thus, each set restricts to one or two group(s), giving approximately num_sets total groups across all sets at the coarsest level. The number of grids needed is determined by the set with the minimum number of groups, since it will be the first to reach a grid with one group. This modifies Equation (21) to be

$$num_g_{min} = \text{floor}\big(\frac{num_groups}{num_sets}\big)\ , \tag{23}$$

$$num_grids = \text{floor}\big(\log_2(num_g_{min})\big) + 2\ . \tag{24}$$

The communication costs and logistical complications of the first strategy seem likely to overwhelm the benefit gained by going to one group instead of num_sets groups. The second strategy was chosen because it involves much less communication and overhead cost. The value of this choice will become apparent in the results section.

With the implemented energy set strategy, there are trade-offs between the number of sets and the number of grids for a fixed number of groups. When there are more sets, more cores can be used at once and wall time should decrease. When there are fewer sets each V-cycle can go deeper, so the preconditioner should be more effective at reducing iteration count.

5. Results

Some tests were done to characterize the impact of preconditioning several problem types, as well as to investigate how different variables affect the preconditioner's behavior. The preconditioning parameters are the Richardson iteration weight, w_k, the number of V-cycles per preconditioner application, and the number of relaxations per level.

The syntax used throughout this section will be that $w\#$ is the weight, $r\#$ is the number of relaxations per level, and $v\#$ is the number of V-cycles, e.g. $w1r1v1$ is one relaxation per level, one V-cycle, and a weight of 1. Using more preconditioning means using larger values of w and/or r and/or v. Other issues investigated are changing the depth of the V-cycle, using a different quadrature set inside the preconditioner than in the rest of the problem, and strong scaling.

The goal of using the preconditioner is to improve convergence behavior of the multigroup Krylov solves. The best metric for measuring this is the total number of multigroup Krylov iterations used in a calculation because it is the most consistent and fair measure. The number of eigenvalue iterations is also compared for

the eigenvalue tests. This is a point of interest rather than a measure of goal attainment. The total number of Krylov iterations is the best proxy for convergence behavior as it encompasses the work that is done within each eigenvalue iteration.

Timing comparisons should be considered heuristically unless otherwise specified. In cases where the calculations were done on a single core, the machine was not dedicated to these calculations and times could vary if the same calculations were repeated. Some problems use an optimized version of the code and others use a debug version. These should give the same iteration count, but not necessarily the same relative times between problems. Further, little effort has been made to optimize the preconditioner for speed. Once the multigrid in energy preconditioner has been optimized for efficiency, the preconditioned times should decrease. How much improvement can be gained is a matter for future study.

Unless otherwise noted, all test problems used a step characteristic spatial solver, level-symmetric angular quadrature, and the grid depth was determined using the default approach, Equation (21) or Equation (24). The Krylov solver was GMRES, which is set to limit the number of multigrid iterations to 1,000 if the problem does not converge earlier. The convergence tolerances are noted for each problem. The tolerance for the multigroup solve is the convergence tolerance used by GMRES in Trilinos [11]. The eigenvalue tolerance is used by PI to determine if the eigenvalue has converged. In Denovo, PI also checks the L2-norm and the infinity-norm of the difference in the fission source between iterations. The default L2-norm tolerance is 1.0 and infinity-norm tolerance is 0.01.

5.1. Fixed Source Parameter Studies

Some fixed source tests are discussed first, with the selection of preconditioning parameters informed by initial scoping tests that are not reported here. Fixed source calculations are particularly useful because the preconditioner can be studied apart from eigenvalue iterations. These tests used a debug version of Denovo on a single core.

The first test was a small, vacuum boundary problem. It contained one material and used ten groups, five upscattering groups, P_0 scattering, an S_4 level-symmetric angular quadrature set, a $3 \times 3 \times 3$ grid, and a multigroup tolerance of 1×10^{-6}. The first three groups contained an isotropic source.

In one set of tests, the weight, w, was varied while the relaxations per level, r, and number of V-cycles, v, were both set to 1 ($w\#r1v1$). The results are in Table 1. In this and all subsequent results an r, v, and/or w of 0 corresponds to the unpreconditioned case. The table heading "Krylov Iters" indicates the number of Krylov iterations needed for convergence.

It is clear that increasing the weight is initially beneficial, reducing the iteration count to 6 from the unpreconditioned 10. After a certain level, increasing weight begins to increase the number of iterations. The number of iterations needed to converge generally continued to increase above a weight of 1.9, and the problem failed to converge beginning with a weight of 2.1. Near a weight of 2 there was some behavior inconsistent with this trend. This issue is under investigation and is not material to the overall conclusions of this work.

Table 1: Small Fixed Source Problem with Vacuum Boundaries, Weight Variation Study

Weight	Relaxations	V-cycles	Krylov Iters
0.0	0	0	10
1.0	1	1	7
1.1	1	1	6
1.2	1	1	6
1.3	1	1	6
1.4	1	1	7
1.5	1	1	8
1.6	1	1	10
1.7	1	1	13
1.8	1	1	20
1.9	1	1	265

Next, the number of relaxations per grid and the number of V-cycles were varied with the weight fixed at 1.0 ($w1r\#v\#$). These results are shown in Table 2. Recall that an r of 3 means the preconditioner performed 3 relaxations per level, and a v of 3 means the preconditioner conducted 3 V-cycles per application.

Initially, increasing r and v reduced the number of iterations needed for convergence. After enough preconditioning was done that only 4 iterations were needed, no additional amount of preconditioning reduced the iteration count further.

Another calculation using a large amount of preconditioning, $w1.3r10v10$, also yielded 4 iterations. This both confirms that more preconditioning did not improve results, and demonstrates that a large amount of preconditioning does not cause breakdown.

The previous problem configuration was modified with reflecting boundary conditions. Both weight and r/v were varied. The multigroup tolerance and was

Table 2: Small Fixed Source Problem with Vacuum Boundaries, Relaxation Count and V-cycle Variation Study

Weight	Relaxations	V-cycles	Krylov Iters
0.0	0	0	10
1.0	1	1	7
1.0	2	2	4
1.0	3	3	4
1.0	4	4	4
1.0	5	5	4
1.0	6	6	4
1.0	7	7	4
1.0	10	10	4

again 1×10^{-6}. Results for varying the weight using 1 relaxation per level and 1 V-cycle are in Table 3. The results for varying the number of relaxations per grid and the number of V-cycles using a weight of 1 can be seen in Table 4. A few r/v variations using a weight of 1.3 instead of a weight of 1 are also in Table 4.

Table 3: Small Fixed Source Problem with Reflecting Boundaries, Weight Variation Study

Weight	Relaxations	V-cycles	Krylov Iters
0.0	0	0	15
1.0	1	1	11
1.1	1	1	10
1.2	1	1	10
1.3	1	1	10
1.4	1	1	11
1.5	1	1	13
1.6	1	1	19
1.7	1	1	50
1.8	1	1	1000

The behavior of this problem was similar to the vacuum case. For this test a weight of 1.8 rather than 2.1 prevented convergence. Like the vacuum test, in-

Table 4: Small Fixed Source Problem with Reflecting Boundaries, Relaxation Count and V-cycle Variation Study

Weight	Relaxations	V-cycles	Krylov Iters
0.0	0	0	15
1.0	1	1	11
1.0	2	2	5
1.0	3	3	3
1.0	4	4	2
1.0	5	5	2
1.0	6	6	2
1.0	7	7	2
1.0	10	10	1
1.3	2	2	4
1.3	3	3	2
1.3	4	4	2
1.3	7	7	2

creasing the weight parameter above 1 decreased the number of iterations initially, and then the number of iterations increased with increasing weight. No anomolous behavior was observed for this problem.

The r and v study showed two new things compared to the vacuum case. One is that with fewer iterations per level and fewer V-cycles, a weight of 1.3 was better than a weight of 1. The other was that the number of iterations could be decreased to 1 when a large amount of preconditioning was done. Whether and when 1 iteration can be reached for a given problem is related to how well Richardson iteration works for that problem's characteristics.

These two sets of results show that increasing r and/or v decreases Krylov iteration count. They also indicate that using a small amount of weight can be beneficial, but a large amount is not. These two small and simple test cases provide a foundation for choosing parameters for testing with larger and more complex problems.

5.2. Fixed Source Angle Comparisons

For preconditioners to be effective, it is helpful if they are fast. One simple way to reduce preconditioning time is to use fewer angles for the preconditioner

solves than for the overall problem. A half iron, half graphite, fixed source problem with an isotropic source assigned to all groups in all spatial cells was used to test this feature. The problem has vacuum boundaries, 27 energy groups, 13 of which have upscattering, and a P_0 scattering expansion. For this calculation a grid of $20 \times 20 \times 20$ was used and the multigroup tolerance was 1×10^{-6}. The preconditioning parameters were $w1r2v2$. The overall problem was solved with an S_8 angular quadrature.

To compare the effect of the reduced angular quadrature, the quadrature order within the preconditioner was varied from S_2 to S_8. The problem was spatially decomposed over four cores and an optimized version of Denovo was used. The results are shown in Table 5.

Table 5: Iron Graphite Fixed Source Problem, Preconditioner Quadrature Order Study

Preconditioner S_N	Time (s)	% Diff Time	Krylov Iters
8	221.66	n/a	17
6	133.93	39.58	17
4	98.67	55.49	17
2	59.15	73.31	18

For this problem, changing the number of solution directions inside the preconditioner had little or no impact on the number of iterations needed for convergence. Reducing the number of ordinates used inside the preconditioner reduced solution time in all cases, with S_2 giving the largest improvement: 73% speedup. As expected, changing the preconditioner did not change the solution more than the convergence tolerance. Using fewer solve directions inside the preconditioner had a very positive impact: it reduced time but did not meaningfully change the solution.

5.3. Multiple Energy Sets

Another important area of investigation was how the preconditioner faired when using multiple energy sets. This preconditioner was built to take advantage of multiple energy sets and be able to scale efficiently.

To investigate the effect of multigrid in energy with multiple energy sets on the Krylov iterations without worrying about the impact on an eigenvalue calculation, the iron-graphite fixed source problem was considered first. These data were calculated using an optimized version of the code.

For this study, the spatial grid was increased to $50 \times 50 \times 50$ and an S_4 quadrature set was used instead of S_8. The unpreconditioned version was compared to one with $w1r2v2$ on 1 to 10 sets. Note that with 27 groups, 10 is the maximum number of energy sets compatible with the preconditioner. The problem was also decomposed over 2×2 spatial blocks. The calculations therefore used between 4 and 40 cores.

The preconditioned calculation took 27 GMRES iterations while the unpreconditioned took 123 iterations, regardless of the number of energy sets used. Because the number of iterations did not change with sets, the only thing to compare is time. The focus is on relative change in time rather than absolute time since there is still room to optimize the preconditioner.

Two plots are shown in Figure 1. The top plot shows the wall time for the preconditioned and unpreconditioned ("regular") calculations as a function of number of energy sets. The bottom plot shows the efficiency of the regular and preconditioned tests, where the base case is 1 set. Figure 2 is a plot of the relative difference between the two times.

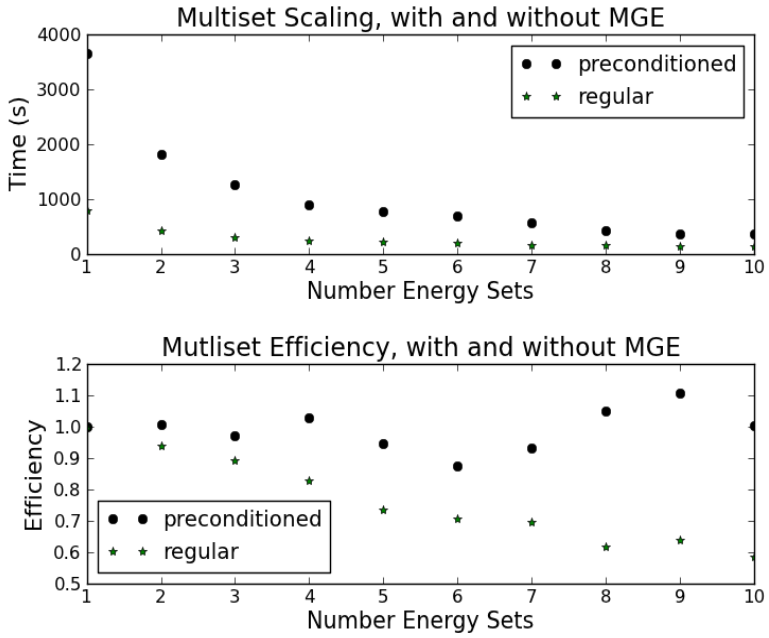

Figure 1: Iron-Graphite Fixed Source Problem, Preconditioned Multiple Energy Set Study

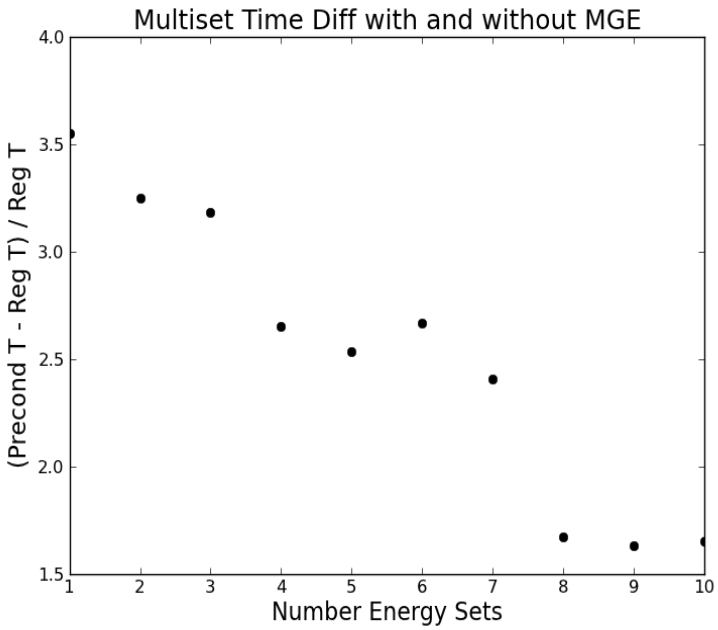

Figure 2: Relative Time Difference In Iron-Graphite Fixed Source Multiple Energy Set Study

The 1-set wall time without preconditioning was 8.00×10^2 seconds and the 10-set time was 1.37×10^2 seconds. Another way to say this is that using 1 set took about 6 times longer than using 10 sets. If the problem scaled linearly in energy, it would have taken 10 times longer. The unpreconditioned efficiency degraded slowly with increasing energy sets, and the efficiency with 10 sets was less than 60%.

With preconditioning, the 1-set time was 3.64×10^3 seconds and the 10-set time was 3.63×10^2 seconds. Going from 1 to 10 sets in this case was linear. The efficiency changed a bit with the number of energy sets, but ranged between about 90% and 110%. Thus, for some sets, the preconditioned case gave better than linear speedup. While the exact efficiencies may not be accurate because of timing variability, the general trend is clear.

The difference between the efficiencies in the preconditioned and regular tests means that the preconditioned tests were accelerated more by using multiple energy sets than the unpreconditioned tests. As a result, the preconditioned times approached the regular times as more sets were used. This can be seen in the

Figure 2 plot of relative difference between the times.

This test was modified and repeated for 1 set and 10 sets. The angular quadrature was S_8 and the problem was not decomposed in space. These choices were made to increase the runtime, ensuring that most of the reported time is solve time rather than setup or completion time. For the preconditioned cases, the preconditioning parameters were increased. Based on the first experiment, it is expected that the scaling will be the same or better when more preconditioning is used.

With $w1.3r4v4$, an unnecessarily-large amount of preconditioning for this problem, but an amount good for investigative purposes, the number of Krylov iterations was 11. With 1 set this took 6.49×10^4 seconds to complete. With 10 sets this was reduced to 5.02×10^3 seconds. A 10-fold increase in computing power gave nearly a 13-fold decrease in run time, or an efficiency of 129%.

Without preconditioning, 124 Krylov iterations were required, the 1-set wall time was 8.81×10^3 seconds, and the 10-set time was 1.04×10^3 seconds. The ratio of 1-set to 10-set time was about 8.5, or 85% efficient. This problem configuration scaled better than the previous configuration whether or not it was preconditioned. The comparison between the preconditioned and unpreconditioned cases show the same trends as the first energy scaling test.

A similar scaling-in-energy test was done with an infinite medium, 27-group, eigenvalue problem. The same trends were observed.

There are some clear reasons for the preconditioner's linear or super-linear scaling in energy. One is that this problem has a group structure that is not always balanced between sets. The multigrid preconditioner mitigates the penalty of energy-group load imbalance. Recall that each set uses the same number of grids, even if the number of groups per set is different. This causes the work in the preconditioner to be energy-load-balanced in all cases. Thus, the relative amount of time spent waiting because of load imbalance decreases when the preconditioner is used.

The main reason for the preconditioner's good scaling is that as the number of sets increases, each application of the preconditioner becomes less costly. The total preconditioning cost goes down because the V-cycle becomes shallower. That is, each application of the preconditioner performs fewer total relaxations, and is therefore less time intensive. This effect becomes more pronounced with larger r and v. If the total number of Krylov iterations remains constant with sets, then there is no trade-off, and the preconditioner simply costs less with increased energy parallelization.

An important outcome of this study is that the number of GMRES iterations did not change with the number of sets in the preconditioned cases. This means

convergence improvement from the preconditioner does not come from the depth of the V-cycle, at least not for the problems tested. Only restricting down one or two grids had as much of an impact as restricting down something like six. This indicates it is not necessary to coarsen to one group.

From an error mode reduction standpoint, this conclusion suggests what a Fourier expansion of the error in energy might look like. Because using a few somewhat-coarse grids had a large impact but using many much coarser grids did not, the bulk of the error might be intermediately oscillatory in energy. If the low-frequency error were dominant, then the coarsest grids would likely be necessary. Only solving on a fine grid, however, was not good enough, so the error is not only high-frequency either. This leaves the presence of modes that are between the two extremes as the likely culprit for slow convergence.

The linear to super-linear energy scaling of the preconditioner warrants more discussion. Restricting down to and relaxing on the coarsest energy grids that were only created in the few-set cases did not provide convergence benefit. Much non-beneficial work was therefore done when only a few sets were used. When many sets were used and only a few grids that were not as coarse were created, all of this work was eliminated without any negative consequences. Thus, the energy scaling was very good.

These results show that limiting the number of energy grids to two or three is likely preferable for many problems since this strategy eliminates the non-beneficial work. By using fewer grids, the preconditioner will probably not cause the entire calculation to scale as well in energy. However, the preconditioner does not communicate between sets and it does load-balance in energy. The preconditioner should then, at worst, leave the energy scaling behavior unchanged.

5.4. V-cycle Depth

The results of the multiple energy set study prompt an investigation of the V-cycle depth. The iron-graphite problem with a $50 \times 50 \times 50$ mesh was used for this purpose. The total problem quadrature was S_8, with S_2 used inside the preconditioner. The preconditioning parameters were $w1r2v2$. The problem was solved with an optimized version of Denovo on 25 cores on the orthanc cluster at ORNL, divided into 5 x-blocks, 5 y-blocks, 1 z-block, and 1 energy set.

The depth of the V-cycle was varied from 1 to 6. A depth of 1 creates only 1 grid and a depth of 6 is the maximum possible for 27 energy groups. The test results are shown in Table 6. The last row is the unpreconditioned case and the time difference is computed as $(t_{unPre} - t)/t_{unPre} * 100$. The minimum number of groups column indicates the minimum number of groups on the coarsest grid.

Table 6: Iron Graphite Fixed Source Problem, V-cycle Depth Study

Depth	Min # Groups	Solve Time (s)	% Diff Time	Krylov Iters
6	1	694.78	53.36	28
5	2	678.78	54.48	28
4	4	643.52	56.80	28
3	7	583.16	60.85	28
2	14	462.77	68.94	28
1	27	387.37	74.00	48
n/a	n/a	1489.69	0.00	124

The 1-grid case was the fastest. Increasing to 2 grids reduced the iteration count, though it also increased the time. Using more than 3 grids increased the solve time but did not change the iteration count. All cases took less time and used fewer Krylov iterations than the unpreconditioned case.

This is the first problem presented in this paper that the preconditioned cases both used fewer iterations and took less time. Part of this success is that the preconditioner took less time per iteration because it used a reduced angular quadrature and a reduced grid depth. It was also the largest and most challenging of the previous test problems.

These results confirm the multiple energy set study findings that using only a few grids is better than using many. The optimal number of grids will be problem dependent, since the shape of the error varies by problem. Without that knowledge *a priori*, choosing a V-cycle depth of 2 is likely to give the best results. A V-cycle depth of 2 provides a significant reduction in iteration count while keeping the solve time low.

5.5. Eigenvalue

All of the test problems presented have been fixed source. The MGE preconditioner will also be useful for eigenvalue calculations because it accelerates the multigroup solves in each eigenvalue iteration. Two eigenvalue test cases using Power Iteration are discussed here. These two tests are the most realistic so far.

The eigenvalue tests have two main goals. One is to see if the preconditioner's general behavior is the same inside of an eigenvalue problem. The other is to see if the reduced Krylov iteration count is a beneficial trade-off for the cost of the preconditioner when used with Power Iteration. Because PI is such a simple

solver, it may not be able to take advantage of the preconditioning as well as some more advanced eigenvalue solvers like Arnoldi or RQI.

5.5.1. 3-D C5G7 Benchmark Study

The preconditioner was applied to the 3-D version of the C5G7 mixed oxide (MOX) benchmark problem [21]. This was solved with an optimized version of Denovo. The problem used three enrichment levels, an external mesh file, 10 materials, 7 groups, 4 upscattering groups, quadruple range angular quadrature using 4 polar and 6 azimuthal angles per octant, a P_0 scattering expansion, a multigroup convergence tolerance (or simply tolerance) of 1×10^{-4}, an eigenvalue tolerance (or k tolerance) of 1×10^{-5}, and k_0 set to 1.14. The problem's dominance ratio is 0.7709, and the reference k is 1.18655 $\pm$ 0.008. The problem was solved on the medium-sized OIC cluster at Oak Ridge, and each calculation was given 720 cores with 40 x-blocks, 18 y-blocks, and 5 z-blocks. The wall time limit was 12 hours.

One drawback of this problem is that it is still fairly simple. The dominance ratio is not particularly high, so PI should not have much difficulty solving it. Further, there are not many energy groups, and the quadrature order is low. The primary goal of this calculation, though, was see whether the lessons learned about preconditioning parameters still hold in an eigenvalue problem. It is also good practice to test new methods on benchmark problems to verify they still calculate the correct solution.

The results are in Table 7. The relative time is found by comparing the preconditioned wall time to the unpreconditioned wall time of 4.46×10^3 seconds. A number larger than 1.0 indicates the preconditioned case took longer. In all cases the computed ks were within the eigenvalue tolerance of one another, but none of them were within the uncertainty bounds of the reported benchmark; they were low by about 0.011. Subsequently, it was determined that using Gauss-Legendre product quarature with 6 polar and 8 azimuthal angles per octant with a linear discontinuous spatial solver instead of QR-4-6 and step characteristic gives the correct k.

The 3-D benchmark study shows the preconditioner with PI can reduce the total number of required Krylov iterations substantially. The number of iterations was cut nearly in half by using $w1r1v1$, and to one-third with $w1.2r1v2$. The $w1.2r1v2$ and $w1.2r2v1$ cases both took less time than $w1r1v1$. These two cases also took less time than the tests using more preconditioning.

An intermediate amount of preconditioning gave the best trade-off between iteration count reduction and computational effort. Using less preconditioning

Table 7: 3-D C5G7 Benchmark, Preconditioning Parameter Scoping with Power Iteration

Weight	Relaxations	V-cycles	Krylov	PI	Rel Time[a]
0.0	0	0	1,224	32	1.00
1.0	1	1	708	32	5.90
1.2	1	2	448	32	5.33
1.2	2	1	448	32	5.37
1.3	2	2	288	32	6.37
1.0	3	3	126	14[b]	9.05
1.5	3	3	192	32	8.36

[a]compared to unpreconditioned, 4.46×10^3 seconds
[b]tolerance = 1×10^{-5}, k tolerance = 1×10^{-3}

cost less per iteration, but the extra iterations cost more time overall. Using more preconditioning reduced the iteration count, but the extra cost per iteration resulted in longer total times.

The number of eigenvalue iterations for a given tolerance set were not changed by preconditioning. This means the preconditioner did not change the solution vector error in a way that had an impact on Power Iteration's behavior.

These results continue to confirm that a small amount of weight works well. Increasing r and v decrease iteration count, but at what can be a high time penalty. The simplicity of the problem, the characteristics of using MGE with PI, or both, keep the preconditioner from being a time win for this test. The main conclusion is that an intermediate amount of preconditioning will likely provide the best balance of reduced iteration count for the time invested for problems that do benefit from the preconditioner.

5.5.2. Full PWR

The next eigenvalue calculation was a whole-core pressurized water reactor (PWR) [7]. The core has a height of 4.2 m, assembly height of 3.6 m, and a lattice pitch of 1.26 cm. There are 157 fuel assemblies and 132 reflector assemblies. Each 17×17 assembly has $1/4$ symmetry and the fuel pins vary from 1.5% to 3.25% enrichment. The pins in the model are homogenized.

In Denovo, each fuel pin has $2 \times 2 \times 100$ spatial cells. The calculations were done using an S_{12} angular quadrature, P_0 scattering, and 44 energy groups, giving a total of 1.73 trillion unknowns. The PWR was decomposed over 137,984 cores

on the Jaguar supercomputer using 578 x- and y-blocks, 700 z-blocks, and 11 energy sets. When the preconditioner was applied, the inside quadrature was S_2, the grid depth was 2, and $w1r2v2$ was selected.

With tolerances of 1×10^{-2}, the unpreconditioned case converged in 24 eigenvalue iterations, using 901 Krylov iterations and taking 92.09 wall-minutes. When preconditioning was added, the eigenvalue iteration count was 28, the number of Krylov iterations was reduced to 312, but the time was increased to 250.09 wall-minutes. The preconditioner reduced the total number of required Krylov iterations by a factor of three, but took much longer. The Denovo-estimated dominance ratio for this tolerance is 0.938.

Both eigenvalue tests demonstrated that the multigrid in energy preconditioner dramatically reduces the Krylov iterations needed for convergence. Power iteration does not take advantage of this benefit. This motivates the need for advanced eigenvalue solvers.

Note that unpreconditioned PI converged these two reactor problems without trouble. We know that there are problems of interest that PI cannot solve quickly, or at all. These are the problems that require advanced eigenvalue solvers like RQI and Arnoldi, and will need preconditioners that can work with those solvers. The MGE preconditioner will likely be more useful for these solvers, and it is the only preconditioner available in Denovo that can operate on the multigroup Krylov solver.

A recent study of RQI in Denovo showed that it did not converge for challenging problems without preconditioning [13]. This is a solver for which the MGE preconditioner could work well. For example, RQI could not solve the 3-D C5G7 benchmark without preconditioning. With preconditioning it behaved very well. There are other initial results showing the MGE preconditioner provides substantial improvement for RQI, and that preconditioned RQI converges more quickly than PI for difficult problems. Details about these results will be reported soon.

6. Conclusions

With larger computers comes the potential to do more challenging neutron transport than ever before. To fulfill that potential, preconditioners that can scale efficiently are required. This paper discussed a new multigrid in energy preconditioner that can scale to hundreds of thousands of cores.

Preconditioning has been an important area of study in the transport community for many years. The use of a multigrid solver as a preconditioner for neutron transport began in the 1980s. Until now, multigrid methods have largely been

conducted in the space and angle dimensions. This is the first multigrid in energy preconditioner for 3-D, large-scale, neutron transport.

A variety of tests were done to characterize the MGE preconditioner implemented in Denovo. These tests showed that performing 1 or 2 relaxations per V-cycle and using 1 or 2 V-cycles per preconditioner application provided the best overall combination of error reduction for the time invested. A weight of 1 when using the weighted Richardson method as the smoother was the most reliable choice.

The preconditioner has the ability to use a lower-order quadrature than the rest of the problem, and this option should be selected whenever practical. Explicit investigation into the effect of limiting V-cycle depth showed that using a fairly shallow depth is sufficient. Truncating the V-cycle reduces calculation time while providing equivalent convergence benefit. The preconditioned cases converged more quickly and in fewer iterations than without preconditioning when these options were combined for the most difficult fixed source test.

The eigenvalue calculations with PI showed that the MGE preconditioner can significantly reduce the total number of Krylov iterations. For these two tests, however, using the preconditioner with PI was not faster than using PI alone. It may be that this preconditioner is not particularly well-suited to PI, or that it will only prove beneficial with more challenging problems. Initial results have shown that the preconditioner works quite well with RQI.

The true potential value of the MGE preconditioner is its ability to capitalize on the new energy scaling capability. The multigroup Krylov solver has been shown to be extremely useful. It converges faster than Gauss Seidel, scales well in energy, and enables calculations with a new level of fidelity. MGE goes hand-in-hand with this solver.

The new preconditioner can operate on entire multigroup vectors unlike other preconditioners in Denovo. MGE inherently provides load balancing across energy groups, and there is no inter-set communication during its application. Thus, it scales well in energy and can scale to hundreds of thousands of cores. It also facilitates the use of very fine energy structures. These are crucial traits for solving "grand challenge" transport problems.

Overall, the MGE preconditioner reduces the number of Krylov iterations required. When parameters and options are chosen to reduce the time spent in the preconditioner, it can provide a time benefit as well. Important areas of future study are solving more difficult problems, using MGE with hundreds of energy groups, and studying its impact on advanced eigenvalue solvers like RQI.

7. Acknowledgements

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. Additional thanks to the Rickover Fellowship Program in Nuclear Engineering sponsored by Naval Reactors Division of the U.S. Department of Energy. This fellowship sponsored the work from which this work is derived.

References

[1] Oak Ridge Leadership Computing Facilities, Jaguar, http://www.olcf.ornl.gov/computing-resources/jaguar/, 10 August 2012.

[2] M. Benzi, Preconditioning Techniques for Large Linear Systems: A Survey, Journal of Computational Physics 182 (2002) 418–477.

[3] T. Evans, G. Davidson, Technical Note, Subject: Parallel Energy Decomposition in Denovo (Rev. 1), Technical Report, Oak Ridge National Laboratory, Oak Ridge, TN, 2010.

[4] E. E. Lewis, W. F. Miller, Jr., Computational Methods of Neutron Transport, American Nuclear Society, La Grange Park, IL, 1 edition, 1993.

[5] M. L. Adams, E. W. Larsen, Fast Iterative Methods for Discrete-Ordinates Particle Transport Calculations, Progress in Nuclear Energy 40 (2002) 3–159.

[6] T. M. Evans, A. S. Stafford, R. N. Slaybaugh, K. T. Clarno, Denovo: A New Three-Dimensional Parallel Discrete Ordinates Code in SCALE, Nuclear Technology 171 (2010) 171–200.

[7] T. M. Evans, G. G. Davidson, R. N. Slaybaugh, Three-Dimensional Full Core Power Calculations for Pressurized Water Reactors, Journal of Physics: Conference Series SciDAC 2010 (2011).

[8] T. M. Evans, Denovo Methods Manual, Technical Report, Oak Ridge National Lab, Oak Ridge, TN, 2009.

[9] R. N. Slaybaugh, Acceleration Methods for Massively Parallel Deterministic Transport, a Ph.D. Dissertation, University of Wisconsin, Madison, WI, 2011.

[10] L. N. Trefethen, D. Bau, III, Numerical Linear Algebra, SIAM, Philadelphia, PA, 1997.

[11] M. A. Heroux, R. A. Bartlett, V. E. Howle, R. J. Hoekstra, J. J. Hu, T. G. Kolda, R. B. Lehoucq, K. R. Long, R. P. Pawlowski, E. T. Phipps, A. G. Salinger, H. K. Thornquist, R. S. Tuminaro, J. M. Willenbring, A. Williams, K. S. Stanley, An overview of the trilinos project, ACM Trans. Math. Softw. 31 (2005) 397–423.

[12] R. S. Baker, K. R. Koch, An S_N Algorithm for the Massively Parallel CM-200 Computer, Nuclear Science and Engineering 128 (1998) 312–320.

[13] R. N. Slaybaugh, T. M. Evans, G. G. Davidson, P. P. H. Wilson, Rayleigh Quotient Iteration in 3D, Deterministic Neutron Transport, in: PHYSOR 2012 Advances in Reactor Physics Linking Research, Industry, and Education, American Nuclear Society, Knoxville, TN, 2012.

[14] D. Knoll, D. E. Keyes, Jacobian-free Newton-Krylov Methods: A Survey of Approaches and Applications, Journal of Computational Physics 193 (2004) 357–397.

[15] S. Oliveira, Y. Deng, Preconditioned Krylov Subspace Methods for Transport Equations, Progress in Nuclear Energy 33 (1998) 155–174.

[16] B. Chang, T. Manteuffel, S. McCormick, J. Ruge, B. Sheehan, Spatial Multigrid for Isotropic Neutron Transport, SIAM Journal on Scientific Computing 29 (2007) 1900–1917.

[17] B. Lee, A Novel Multigrid Method for S_N Discretizations of the Mono-Energetic Boltzmann Transport Equation in the Optically Thick and Thin Regimes With Anisotropic Scattering, Part I, SIAM Journal on Scientific Computing 31 (2010) 4744–4773.

[18] B. T. Adams, J. E. Morel, A Two-Grid Acceleration Scheme for the Multigroup S_N Equations With Neutron Upscattering, Nuclear Science and Engineering 115 (1993) 253–264.

[19] W. L. Briggs, V. E. Henson, S. F. McCormick, A Multigrid Tutorial, SIAM, Philadelphia, PA, second edition, 2000.

[20] WolframMathWorld, Binary Tree, http://mathworld.wolfram.com/BinaryTree.html, 14 September 2012.

[21] OECD-NEA, Benchmark on Deterministic Transport Calculations Without Spatial Homogenisation, MOX Fuel Assembly 3-D Extension Case, Technical Report 5420, Nuclear Energy Agency, Organisation for Economic Cooperation and Development, Paris, France, 2005.